Recent Advances in Fibre Orientation Estimation Using Diffusion-Weighted MRI

J-Donald Tournier
Brain Research Institute, Melbourne, Australia

INTRODUCTION

Diffusion-weighted (DW) MRI is sensitive to the microscopic thermal motion of water molecules on scales comparable with cellular structures. In biological tissue, the motion of these water molecules is hindered by a range of obstacles such as cell membranes, making it possible to probe microscopic properties of the tissue non-invasively in-vivo. Due to the coherent arrangement of fibrous structures in ordered tissue such as brain white matter or muscle, there is a dependence of the measured DW signal on the orientation along which the measurement is performed, due to the larger degree of hindrance to diffusion in directions perpendicular to the orientation of the fibres. Therefore, in principle DW-MRI can be used to infer the orientation of white matter fibre tracts in each voxel in the in-vivo brain. This information can potentially be used to track white matter fascicles through the brain to establish maps of connectivity between the different brain regions [1].

The diffusion tensor model was first proposed to describe the orientation dependence of the DW signal [2]. It is currently the most widely used framework for charaterising DW data, and has been used in a wide range of applications. The principal orientation of the diffusion tensor (its major eigenvector) has been shown to correspond well with the direction of the fibres in major white matter tracts, such as the optic radiations [3]. Based on this premise, a large number of fibre-tracking algorithms have been proposed that rely on the diffusion tensor model to provide a reliable estimate of the orientation of white matter fibres (e.g. [1,4-8]).

However, it is becoming increasingly clear that the diffusion tensor model is inadequate in regions containing multiple distinct fibre orientations [9-11]. There are two main reasons for this. First, the diffusion tensor model is based on the assumption of free diffusion, which is clearly not the case in biological tissues. Second, and more importantly, the diffusion tensor can only possess a single orientation maximum, and is therefore incapable of characterising a multi-orientation system. Unfortunately, there are many regions containing multiple fibre orientations in the brain, and many tracts will be affected by this problem at some point along their path (such a region is shown in 1). Given that the orientations provided by the diffusion tensor are unreliable in such regions, this problem can have a considerable knock-on effect on the reliability of fibre-tracking algorithms.

This paper considers the effects of the so-called 'crossing fibre' problem on DW-MRI, and reviews some of the techniques that have been proposed so far to estimate fibre orientations in such cases.

THE CROSSING FIBRE PROBLEM

DW-MRI acquisitions typically have low resolution, due to the SNR and scan time limitations of the technique. This leads to considerable problems with partial volume effects, as the signal measured in the voxel may contain contributions from different fibre tracts. This situation will arise in voxels located at the interface between two different fibre bundles that brush past each other without actually crossing. A range of orientations will also be found in areas where fibres diverge or 'fan out', or simply where the curvature of the fibre tract is comparable with the voxel size. Thus, it is not necessary for fibres from the different tracts to actually intermesh for a 'crossing fibre' problem to occur: any partial volume effects between tracts with different orientations, or even within a single tract with a range of orientations, cannot be adequately described by the diffusion tensor model.

Impact on the DW signal

In 'crossing fibre' voxels, the DW signal will contain contributions from each of the various fibre bundles present. In the slow exchange limit, these contributions will add up linearly. In other words, as long as the number of water molecules exchanging between fibre bundles over the timescale of the experiment is negligible, then the measured DW signal will correspond to the linear sum of the DW signal originating from each of the individual components. The diffusion time in a typical MRI experiment is on the order of 50-100 ms, during which time water molecules will move on average by 5 to 10 μm in normal brain tissue. It is thus reasonable to assume a slow exchange regime as long as fibres from the different bundles do not intermesh on a scale of less than a few tens of microns. This may be justified in normal white matter, where axons are usually bundled into fascicles whose diameter is typically greater than a few hundred microns. However, given that the diameter of a single neuronal axon ranges typically from 1 to 10 μm, it should be noted that there is potential for fibres to interdigitate on scales that may invalidate the slow exchange regime assumption, which would complicate the interpretation of DW-MRI data further.

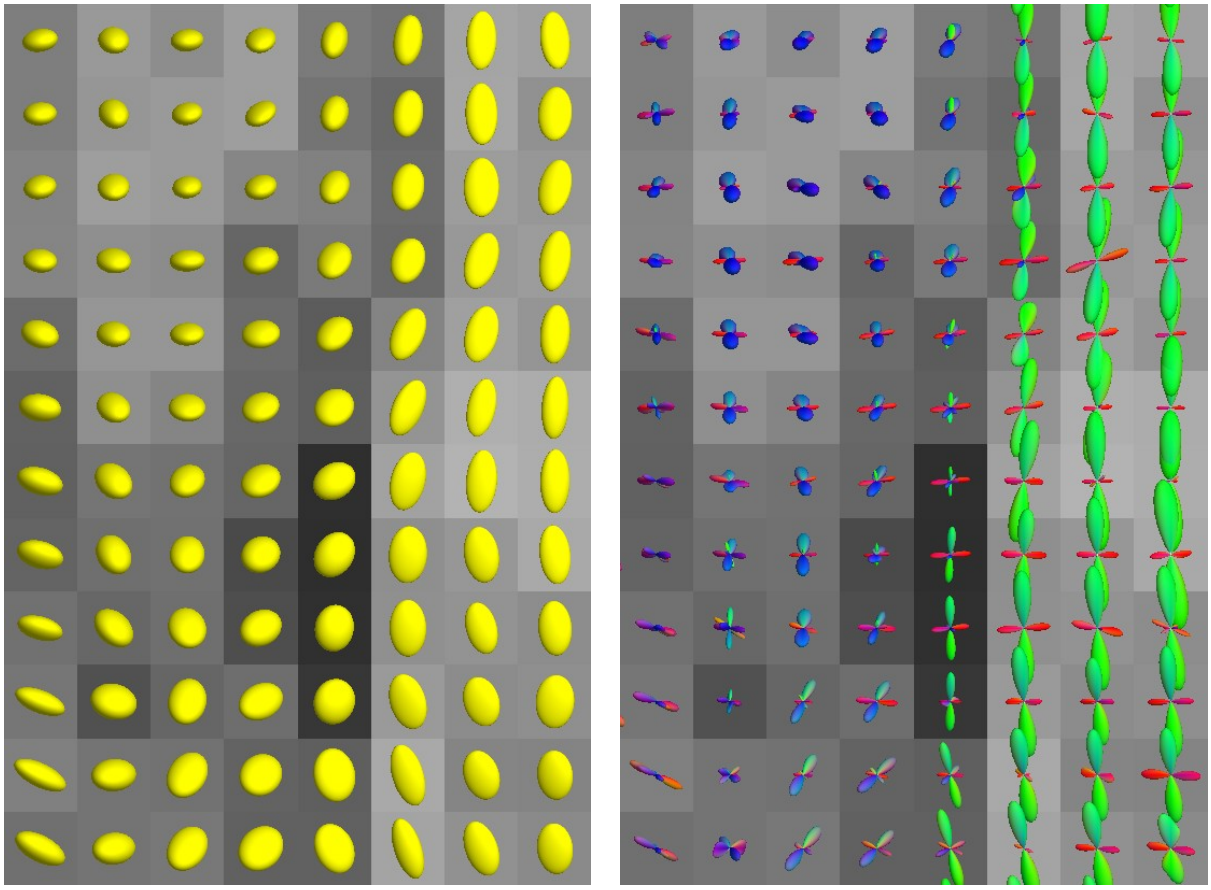

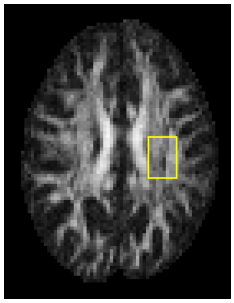

Figure 1: an example of crossing fibres in the human brain. The ROI is located in the supraventricular white matter, as shown in the axial FA map on the left. Top left: results obtained using the diffusion tensor model. Top right, results obtained using the spherical deconvolution (described here), illustrating the extent of fibre crossing in the region. As can be appreciated, the diffusion tensor model is unable to detect or characterise the multiple fibre orientations present. The fibre orientations picked out by the spherical deconvolution correspond to the projections from the corpus callosum running left-right (red), the fibres of the corona radiata running through the plane (blue), and the arcuate fasciculus running anterior-posterior (green). Data acquired on a 1.5T Siemens Avanto, $2.1 \times 2.1 \times 3$ mm^3 voxels, 37 contiguous slices, 60 DW directions, $b = 3{,}000$ s/mm^2.

This concept can also be applied to curving fibres: provided the radius of curvature of the fibres remains small compared to the diffusion distance, the fibres can be considered to be linear over the length scales probed by individual spins. There is effectively no exchange between portions of the fibre tract that have significantly different orientations. This can also be extended to diverging fibres: as long as the divergence occurs on length scales much greater than the diffusion distance, it can be assumed that there is no exchange between differently oriented portions of the same fibre tract.

Orientation information representation

A further issue to consider is how the orientation information is represented. Certain algorithms attempt to estimate the number and orientations of a set of discrete fibre bundles, whereas others will provide an estimate of the distribution of fibre orientations present within the voxel. The discrete representation works well when the number of fibre bundles is known *a priori*, and when there is negligible spread about each orientation. Also, for small numbers of orientations (the typical case), few parameters are required to characterise the system, leading in theory to better conditioned reconstructions. Unfortunately, this typically relies on independent techniques for estimating the number of orientations present, which may themselves introduce errors into the reconstruction. Also, any significant curvature or divergence within a voxel would introduce a range or orientations that would not be well characterised by these approaches. Finally, in practice most methods of this type limit themselves to two orientations due to numerical instability.

On the other hand, the continuous fibre orientation distribution (FOD) representation is much more flexible and

in theory can be used to characterise any fibre arrangement. Unfortunately, the number of parameters required to characterise the distribution faithfully is large, leading to poorly-conditioned reconstructions. To combat this problem, it is common to use filtering approaches, which typically have a negative impact on the angular resolution of the distribution. This can lead to an overestimation of the amount of spread about each of the orientations, and can reduce the ability of the technique to resolve fibre orientations that are separated by small angles.

Data acquisition strategies

In the context of fibre orientation estimation, the ultimate aim of acquiring DW data will typically be to perform fibre-tracking experiments that require whole brain coverage. In addition, the number of DW orientations typically required to resolve multiple fibre orientations is much greater than for standard DTI. This is part of the reason why most imaging is done using fast multi-slice DW EPI sequences.

For the DW orientation sampling scheme, there two main acquisition strategies. The most common and practical approach is to sample a relatively large number of orientations evenly distributed over spherical half-space with a constant *b*-value. Sampling schemes containing from 40 to over 200 orientations have been used in the literature. Such approaches are collectively known as high angular resolution DW imaging (HARDI). This is the favoured approach for routine human in vivo studies, since it generally allows much shorter scan times than the alternatives.

The optimal *b*-value (i.e. the strength of the diffusion encoding) to use with HARDI sequences is yet to be fully determined. However, it is clear that fibre crossings can be more readily detected and characterised if *b*-values greater than approximately 2,000 s/mm^2 are used, rather than the b=1,000 s/mm^2 typically used for standard DTI sequences [12]. Higher *b*-values lead to lower SNR in the DW images, and also require stronger gradient hardware. Therefore, the *b*-value used in practice is a compromise between SNR, gradient strength, and the performance of the reconstruction algorithm.

Another strategy is to apply each diffusion weighting according to a Cartesian grid in *q*-space, producing data that are then suitable for *q*-space type analyses (see below for details). Unfortunately, this leads to impractical scan times, as the number of distinct diffusion weightings required is of the order of 500. Currently, the only technique to use this type of acquisition for estimating fibre orientations is diffusion spectrum imaging (DSI), which is also discussed below.

THE *Q*-SPACE FORMALISM

A number of techniques discussed here are based on the *q*-space framework. This formalism introduces the concept of a spin propagator $P(\mathbf{r}_0, \mathbf{r}, \Delta)$, giving the probability of a spin initially at position $\mathbf{r}_0$ diffusing by a distance $\mathbf{r}$ during the diffusion time Δ. Under certain conditions (see below), it can be shown that the DW signal attenuation $E(\mathbf{r}_0, \mathbf{q}, \Delta)$ is related to the spin propagator via a simple Fourier relationship:

$$E(\mathbf{r_0}, \mathbf{q}, \Delta) = \int P(\mathbf{r_0}, \mathbf{r}, \Delta)\ e^{i 2\pi \mathbf{q}\cdot\mathbf{r}}\ \mathrm{d}\mathbf{r} \qquad (1)$$

where $\mathbf{q}$ is the reciprocal wave-vector, defined as $\gamma\delta\mathbf{g}/2\pi$. The simplest type of experiment that can be performed using this approach involves acquiring signal samples at regular intervals in *q*-space while keeping Δ constant, so that a simple fast Fourier transform yields the spin propagator, averaged over the imaging voxel. Note that the *q*-space approach does not assume Gaussian free diffusion, as is the case in traditional diffusion tensor imaging, and therefore provides a much more complete and flexible representation of the diffusion process.

Eq (1) holds when the narrow pulse approximation is valid, i.e. when the DW gradient pulse duration δ is negligible compared to the diffusion time Δ (the delay between the two DW gradient pulses). Unfortunately, this assumption cannot be met in practice on clinical systems due to the huge demands on the gradient hardware. In this case, the DW gradient pulse duration δ is finite, and $\mathbf{r}$ (the distance travelled during the diffusion time) is poorly defined. This can be accounted for with a simple modification of eq (1), whereby $\mathbf{r}$ now corresponds to the distance between the position of the spin averaged over the duration of the first DW gradient pulse, to the corresponding position averaged over the second gradient pulse duration [13].

Although the spin propagator provides the most complete description of the diffusion process, its relationship to the fibre orientation distribution remains to be fully elucidated. It is widely assumed that 'ridges' (i.e. large concentrations along a particular orientation) in the propagator must correspond to fibre orientations. However, using the maxima in the propagator as fibre orientation estimates forces the use of the discrete representation for the orientation information, with all its associated disadvantages as discussed previously.

RECONSTRUCTION ALGORITHMS

A large number of methods have been proposed to estimate fibre orientations in the presence of crossing fibres. These can be categorised into three main groups: those building on the diffusion tensor model, those relying on the *q*-space framework, and those based on empirical observation. Each of these methods is discussed here, according to their categorisation.

Tensor based methods

Multiple tensor fitting

This approach relies on the fact that the diffusion tensor model provides a good fit to the observed DW data when a single fibre population is present. It is therefore reasonable to assert that the DW signal in voxels containing several distinct fibre bundles is well characterised using a mixture of tensors [11,14,15]. In other words:

$$S_i = S_0 \sum_{j=0}^{N} f_j \; e^{-\boldsymbol{b}_i : \boldsymbol{D}_j} \tag{2}$$

where S_i is the DW signal measured along direction i, S_0 is the signal measured without any diffusion weighting, N is the assumed number of fibre populations, f_j is the volume fraction of the jth fibre population, $\boldsymbol{b}_i$ is the b-matrix associated with direction i, and $\boldsymbol{D}_j$ is the diffusion tensor corresponding to the jth fibre population. The different approaches proposed to date differ mostly in the method used to solve eq (2). An independent method is typically required to estimate the number of fibre populations present, often using a statistical test to determine which model provides the best fit the data. Most implementations also apply some additional constraints on each component diffusion tensor to reduce the number of degrees of freedom, either by fixing the eigenvalues [11], or by restricting the tensor to be axially symmetric (i.e. $\lambda_2 = \lambda_3$) [14,15]. Despite this, the maximum number of fibre populations is typically restricted to $N = 2$, since higher numbers tend to yield unstable solutions. Also, these approaches suffer from the limitations associated with the use of a discrete orientation representation, as discussed previously.

Fibre Orientation Estimated using Continuous Axially Symmetric Tensors (FORECAST)

The multiple tensor fitting approach can be generalised by letting N in eq (2) tend to infinity [16]. This leads to an equation of the form:

$$S(\theta,\phi) = S_0 \int_0^{2\pi} \int_0^{\pi} F(\theta',\phi') \; e^{\boldsymbol{b}(\theta,\phi):\boldsymbol{D}(\theta,\phi)} \; \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi \tag{3}$$

where $S(\theta,\phi)$ is the DW signal measured by applying the DW gradients along the direction (θ,ϕ), $F(\theta,\phi)$ is the unknown fibre orientation distribution (FOD), $\boldsymbol{b}(\theta,\phi)$ is the b-matrix obtained by applying the DW gradients along the direction (θ,ϕ), and $\boldsymbol{D}(\theta,\phi)$ is the diffusion tensor corresponding to a fibre bundle aligned along (θ,ϕ). For this problem to be tractable, the eigenvalues of $\boldsymbol{D}(\theta,\phi)$ need to fixed; in other words the shape of the diffusion tensor is assumed to be the same for all fibre bundles within the voxel. In Anderson's implementation, the major to minor eigenvalue ratio is determined from the data themselves, assuming axially symmetric tensors [16]. Eq (3) can then be solved by expressing the problem using a suitable set of basis functions, such as spherical harmonics. In contrast to many of the other approaches, this technique provides a direct estimate of the fibre orientation distribution. Finally, as this method uses the continuous representation, it has the advantages and disadvantages associated with it (see the section 'Orientation information representation' above).

q-space based methods

Diffusion Spectrum Imaging (DSI)

This method is based directly on the *q*-space formalism, and provides an estimate of the average spin propagator for each imaging voxel by acquiring images with **q** vectors arranged in a rectilinear 10×10×10 grid, suitable for direct Fourier transformation [17]. Thanks to the antipodal symmetry of the diffusion process, the number of images needed can be halved, giving a total of 500 images that need to be acquired, each with their own distinct diffusion-encoding **q** vector. Also, since the DW signal can be shown to be theoretically real and positive, the Fourier transform can be performed using the signal modulus, rather than its complex representation, thus bypassing problems associated with motion-induced phase errors.

In this way, an estimate of the 3D spin propagator can be obtained for each voxel. Since it is the orientation information that is of interest, the radial part of the propagator is typically projected out by performing a weighted radial summation to yield the so-called orientation distribution function (ODF):

$$\mathrm{ODF}(\hat{\mathbf{u}}) = \int_0^\infty P(\rho\,\hat{\mathbf{u}})\rho^2 \, \mathrm{d}\rho \tag{4}$$

where $\hat{\mathbf{u}}$ is a unit vector, and $P(\mathbf{r})$ is the spin propagator for the voxel of interest. It is then assumed that maxima in the ODF correspond to the orientations of the fibre bundles.

Although DSI provides the most complete characterisation of the diffusion process within each voxel, there are some distinct disadvantages to its use. First, the large number of images required to reconstruct the spin propagator translates into prohibitively long scan times. Second, to avoid truncation artefacts, *q*-space needs to be sampled over a large range, placing huge demands on the gradient hardware. In Wedeen's implementation, the largest *q*-value used corresponds to b_{max} = 17,000 s/mm^2, compared to b_{max} = 1,000 s/mm^2 typically used for DTI data acquisitions [17]. Third, a large amount of the data collected is actually thrown away by the radial projection (this concern is largely addressed with *Q*-Ball imaging – see below). Finally, the ODF does not correspond to the fibre orientation distribution, complicating its interpretation. As mentioned above for *q*-space, since peaks are interpreted as discrete fibre orientations, the technique has all the limitations associated with the discrete representation.

Q-Ball Imaging (QBI)

As mentioned above, DSI is relatively inefficient in its sampling strategy, due to the radial projection of the data. *Q*-ball imaging (QBI) provides a good approximation to an ODF similar to that produced by DSI, using the so-called Funk-Radon transform (FRT) [18]. The FRT provides a mapping between two functions defined over spherical coordinates. The FRT $F'(\theta',\phi')$ of a function $F(\theta,\phi)$ is defined as the integral of $F(\theta,\phi)$ performed over the equator that lies in the plane perpendicular to the orientation (θ',ϕ'). It can be shown that the FRT of a spherical shell in *q*-space is closely related to the radial projection of the spin propagator: the radial integral is now performed over a *beam* in *q*-space, rather than along an infinitely narrow line. The cross-section of this beam is given by a zeroth-order Bessel function, the width of which is inversely proportional to the radius of the spherical *q*-space shell over which the FRT was performed. This technique is thus appropriate for HARDI data, and its resolution is related to the *b*-value used. In practice, the ODF produced by QBI has relatively poor angular contrast, with small amplitude differences between peaks and troughs. The ODF is therefore scaled for display purposes so that its minimum and maximum values are mapped to 0 and 1 respectively. This improves the visualisation of the ODF, and multiple fibre orientations can readily be observed.

QBI has been performed with *b*-values ranging from 4,000 s/mm^2 to 12,000 s/mm^2 [19]; in practice the optimal value is a compromise between SNR and angular resolution. With this technique it is possible to obtain information similar to that provided by DSI with much shorter scan times. However, the ODF produced by QBI is subject to the same limitations as that provided by DSI; in particular the QBI ODF is also not equivalent of the fibre orientation distribution.

Persistent Angular Structure MRI (PAS-MRI)

This method defines a function called the persistent angular structure (PAS) $p(\hat{\mathbf{u}})$, which is related to the spin propagator $P(\mathbf{u})$ by the equation $P(\mathbf{u}) = p(\hat{\mathbf{u}})\ r^{-2}\ \delta(|\mathbf{u}| - r)$ [20]. The assumption here is that the spin propagator has zero amplitude everywhere but on a spherical shell; in other words that all spins in the system travel by the same distance during the diffusion time. Therefore, the PAS can be thought of as a projection of the angular component of the spin propagator, ignoring any radial information (akin to the ODF in DSI and QBI). Using this simplified spin propagator, it becomes possible to perform the *q*-space Fourier transform in eq (1) using HARDI data, acquired on a shell in *q*-space. PAS-MRI also includes a maximum entropy constraint to better condition the problem. The estimated PAS therefore contains the minimum amount of information required for it to remain consistent with the data.

The equation to be solved is non-linear, and requires a computer-intensive iterative minimisation procedure. PAS-MRI use the Levenberg–Marquardt algorithm, with a Gaussian noise model. This is one of the main limitation of the method, as reconstruction times can run into days for a full 3D data set. Also, as with DSI and QBI, the PAS does not correspond directly to the fibre orientation distribution, and the relationship between them is not clear; peaks in the PAS are assumed to provide the fibre orientations. A further concern is that the maximum entropy constraint may introduce artefacts into the PAS, for example by suppressing smaller peaks or by merging two fibre orientations that are separated by a small angle into a single peak. However, PAS-MRI is able to extract sensible information from relatively low *b*-value HARDI data, which would be unsuitable for use with other reconstruction techniques. Moreover, the angular resolution achievable with PAS-MRI is very high due to the maximum entropy constraint, especially when compared to QBI or DSI.

Composite Hindered And Restricted Model of Diffusion (CHARMED)

With this technique, the signal is assumed to originate from a set of non-exchanging compartments, each of which can be approximated by a simple model [21]. Spins outside the fibres are assumed to diffuse in a hindered

environment that is well approximated by an axially symmetric diffusion tensor, whereas spins inside the fibres are assumed to be fully restricted. The intra-fibre compartment can be decomposed into a component parallel to the fibre orientation, modelled as free Gaussian diffusion, and a perpendicular component that can be modelled using an analytical expression for diffusion within a cylinder [21]. The voxel under consideration can then be modelled using any combination of the different compartments. In practice it was found that the combination of one hindered and two restricted compartments provides the best fit in the presence of two fibre populations.

In [21], data were acquired from excised spinal cord at multiple *b*-values up to b_{max} = 44,000 s/mm^2, along 31 directions, and the total number of images was 496. Due to the difficulty of acquiring this kind of data on a clinical system, it is not clear whether the CHARMED model can be used for human in vivo studies. Also, the orientation information is represented using a set of discrete orientations, which implies some of the limitations discussed previously.

Methods based on empirical observation

Spherical deconvolution

Spherical deconvolution uses HARDI data, and is based on two assumptions. First, the exchange between fibre bundles with distinct orientations is slow (see the section 'Impact on the DW signal' above). Second, the DW signal profile (i.e. the DW signal measured as a function of orientation, using the fibre tract's frame of reference) is constant for all white matter fibre bundles in the white matter, and can be represented using an axially symmetric *response function* $R(\theta)$. The DW signal $S(\theta,\varphi)$ that is actually measured during a HARDI experiment is then given by the convolution over spherical coordinates of the response function $R(\theta)$ with $F(\theta,\varphi)$, the fibre orientation distribution (FOD). Therefore, the FOD can be obtained by performing the spherical deconvolution of $R(\theta)$ from $S(\theta,\varphi)$. The response function $R(\theta)$ can be estimated from the data themselves, by using the DW signal from a set of voxels assumed to contain a single coherently oriented fibre population. The method is therefore independent of any model of diffusion, relying purely on the observed DW signal.

The spherical deconvolution operation can be performed efficiently using the spherical harmonic framework [22]. However, the operation is sensitive to noise, and some form of regularisation is needed. A method has been proposed to find an optimal low-pass filter, which produces robust results at the expense of angular resolution [23]. A more promising approach is to constrain negative values in the FOD, since they are clearly unphysical. This results in significant improvements in the angular resolution achievable with spherical deconvolution, as illustrated in figure 1.

Spherical deconvolution is based on the observation that the DW signal profile does not vary significantly between different normal white matter fibre bundles [22]. The DW signal is always lowest along the direction of the fibres, and largest in the perpendicular plane. The exact shape is dependent on the *b*-value, and specific properties of the fibre bundle, such as the average axonal diameter, axonal density and the thickness of the myelin. Fortunately, when averaged over the scale of a voxel, these parameters vary relatively little across the white matter in the normal brain, and the impact of these variations on the DW signal profile remains moderate. Moreover, although this effect may introduce errors in the estimated volume fractions, the orientations are not affected [22]. Finally, higher *b*-values introduce stronger suppression of the DW signal measured parallel to the fibres than in the perpendicular plane, so that the DW signal profile assumes a 'flatter' shape that should be less dependent on the properties of the fibre bundle. A further benefit of using high *b*-values is an increase in the theoretically achievable angular resolution, due to the increased angular contrast in the DW signal, although in practice this has to be balanced against SNR.

CONCLUSION

As has been shown here, the single diffusion tensor model is often unable to extract the desired fibre orientations from DW-MRI data. This has a negative impact on fibre-tracking techniques that rely on it to provide a good estimate of the fibre orientation, as many fibre tracts in the brain will encounter fibre crossing regions at some point, including some of the better known pathways such as the corticospinal tracts. There is therefore an increasing recognition that techniques capable of resolving multiple fibre orientations within each voxel are essential for reliable fibre tracking using DW-MRI to be achieved. This paper has reviewed the most promising techniques for dealing with this issue. Many of these are now beginning to be used in fibre tracking algorithms, and will eventually lead to tractography methods that can robustly track through crossing fibre regions.

REFERENCES

[1] Mori S, Crain BJ, Chacko VP, van Zijl PCM. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann. Neurol. 1999; **45:**265-269.

[2] Basser PJ, Mattiello J, Le Bihan D. Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson. B 1994; **103:**247-254.

[3] Lin CP, Tseng WY, Cheng HC, Chen JH. Validation of diffusion tensor magnetic resonance axonal fiber imaging with registered manganese-enhanced optic tracts. NeuroImage 2001; **14:**1035-1047.

[4] Conturo TE, Lori NF, Cull TS, Akbudak E, Snyder AZ, Shimony JS, McKinstry RC, Burton H, Raichle ME. Tracking neuronal fiber pathways in the living human brain. Proc. Nat. Acad. Sci. USA 1999; **96:**10422-10427.

[5] Parker GJM, Wheeler-Kingshott CAM, Barker GJ. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. IEEE Trans. Med. Imag. 2002; **21:**505-512.

[6] Parker GJM, Haroon HA, Wheeler-Kingshott CAM. A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. J. Magn. Reson. Imag. 2003; **18:**242-254.

[7] Behrens TEJ, Johanse-Berg H, Woolrich MW, Smith SM, Wheeler-Kingshott CAM, Boulby PA, Barker GJ, Sillery EL, Sheehan K, Ciccarelli O, Thompson AJ, Brady JM, Matthews PM. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. Nat. Neurosci. 2003; **6:**750-757.

[8] Tournier JD, Calamante F, Gadian DG, Connelly A. Diffusion-weighted MRI fibre-tracking using a front evolution algorithm. NeuroImage 2003; **20:**276-288.

[9] Alexander DC, Barker GJ, Arridge SR. Detection and modelling of non-Gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 2002; **48:**331-340.

[10] Frank LR. Characterisation of anisotropy in high angular resolution diffusion-weighted MRI. Magn. Reson. Med. 2002; **47:**1083-1099.

[11] Tuch DS, Reese TG, Wiegell MR, Makris N, Belliveau JW, Wedeen VJ. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magn. Reson. Med. 2002; **48:**577-582.

[12] Alexander AL, Hasan KM, Lazar M, Tsuruda JS, Parker DL. Analysis of partial volume effects in diffusion-tensor MRI. Magn. Reson. Med. 2001; **45:**770-780.

[13] Mitra PP, Halperin BI. Effects of finite gradient-pulse widths in pulsed-gradient diffusion measurements. J. Magn. Reson. A 1995; **113:**94-101.

[14] Hosey T, Williams G, Ansorge R. Inference of multiple fiber orientations in high angular resolution diffusion imaging. Magn. Reson. Med. 2005; **54:**1480-1489.

[15] Kreher BW, Schneider JF, Mader I, Martin E, Hennig J, Il'yasov KA. Multitensor approach for analysis and tracking of complex fiber configurations. Magn. Reson. Med. 2005; **54:**1216-1225.

[16] Anderson AW. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. Magn. Reson. Med. 2005; **54:**1194-1206.

[17] Wedeen VJ, Hagmann P, Tseng W-YI, Reese TG, Weisskoff RM. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn. Reson. Med. 2005; **54:**1377-1386.

[18] Tuch DS. Q-ball imaging. Magn. Reson. Med. 2004; **52:**1358-1372.

[19] Tuch DS, Reese TG, Wiegell MR, Wedeen VJ. Q-ball imaging. Proceedings of the 11th annual meeting of the ISMRM, Toronto, Canada. 2003; p. 63.

[20] Jansons KM, Alexander DC. Persistent angular structure: new insights from diffusion MRI data. Inverse Problems 2003; **19:**1031-1046.

[21] Assaf Y, Freidlin RZ, Rohde GK, Basser PJ. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. Magn. Reson. Med. 2004; **52:**965-978.

[22] Tournier J-D, Calamante F, Gadian DG, Connelly A. Direct estimation of the fibre orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage 2004; **23:**1176-1185.

[23] Tournier J-D, Calamante F, Connelly A. Improved characterisation of crossing fibres: optimisation of spherical deconvolution parameters using a minimum entropy principle. Proceedings of the 13th annual meeting of the ISMRM, Miami, USA. 2005; p 384.